

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES**

Application No.: 10/813,857

Filed: March 31, 2004

Inventor(s):

Robert E. Cypher

§ Examiner: Eland, Shawn
§ Group/Art Unit: 2188
§ Atty. Dkt. No: 5681-13501
§ Confirm No. 5571

Title: Multi-Node System with Response Information in Memory

APPEAL BRIEF

Mail Stop Appeal Brief - Patents

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450

Sir/Madam:

Further to the Notice of Appeal of October 1, 2007, Appellant presents this Appeal Brief. Appellant respectfully requests that this appeal be considered by the Board of Patent Appeals and Interferences.

I. REAL PARTY IN INTEREST

The subject application is owned by Sun Microsystems, Inc. An assignment of the present application to the owner is recorded at Reel 015164, Frame 0486.

II. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences known to Appellant.

III. STATUS OF CLAIMS

Claims 1-48 are pending. Claims 1-48 are rejected under 35 U.S.C. § 102(b). It is these rejections that are being appealed. A copy of claims 1-48 is included in the Claims Appendix attached hereto.

IV. STATUS OF AMENDMEMNTS

No unentered amendment to the claims has been filed after final rejection. The Appendix hereto reflects the current state of the rejected claims.

V. SUMMARY OF THE CLAIMED SUBJECT MATTER

Independent claim 1 is directed to a system (*See e.g.*, FIG. 20, #100) including a plurality of nodes (*See e.g.*, FIG. 20, #140A-C) coupled by an inter-node network (*See e.g.*, FIG. 20, #154). Each node includes a plurality of active devices (*See e.g.*, FIG. 1 and FIG. 20, #142, #146, and specification page 11, lines 16-27), a memory subsystem (*See e.g.*, FIG. 20, #144), and an address network (*See e.g.*, FIG. 1 and FIG. 20, #150) and a data network (*See e.g.*, FIG. 1 and FIG. 20, #152) respectively configured to convey address packets and data packets between the active devices and the memory subsystem (*See e.g.*, specification page 11, lines 16-27 and page 62, lines 5-21). The memory subsystem included in a node of the plurality of nodes is configured to maintain a response indication (*See e.g.*, FIG. 35) indicating whether the memory subsystem should

send a data packet corresponding to a coherency unit in response to receiving an address packet requesting an access right to the coherency unit from an active device in the same node (*See, e.g.*, specification page 103, lines 18-23). The node is also configured to store a node identifier for the coherency unit. The node identifier identifies a different node of the plurality nodes in which the coherency unit is in a modified global access state (*See, e.g.*, specification page 67, line 27-page 68, line 1).

Independent claim 17 is directed to a node *See e.g., FIG. 20, #140A-C*) for use in a multi-node computer system (*See e.g., FIG. 20, #100*). The node includes a plurality of client devices including a memory subsystem (*See e.g., FIG. 20, #144*), an active device (*See e.g., FIG. 20, #142, #146*), and an interface (*See e.g., FIG. 20, #148*) configured to send and receive coherency messages on an inter-node network coupling nodes in the multi-node computer system (*See e.g., FIG. 20, #154* and specification page 63, lines 7-10). The node also includes an address network configured to convey address packets between the plurality of client devices (*See e.g., FIG. 1 and FIG. 20, #150*), and a data network configured to convey data packets between the plurality of client devices (*See e.g., FIG. 1 and FIG. 20, #152* and specification page 11, lines 16-27 and page 62, lines 5-21). The memory may be configured to maintain a response indication (*See e.g., FIG. 35*) indicating whether the memory should send a data packet corresponding to a coherency unit on the data network in response to receiving an address packet requesting an access right to the coherency unit from the active device (*See, e.g., specification page 103, lines 18-23*). The node may be further configured to store a node identifier for the coherency unit, where the node identifier identifies a different node in the multi-node system in which the coherency unit is in a modified global access state (*See, e.g., specification page 67, line 27-page 68, line 1*).

Independent claim 33 is directed to a method of operating a multi-node computer system (*See e.g., FIG. 20, #100*) that includes a plurality of nodes (*See e.g., FIG. 20, #140A-C*) coupled by an inter-node network *See e.g., FIG. 20, #154*). Each node of the system includes an active device (*See e.g., FIG. 20, #142, #146*), a memory subsystem

(*See* e.g., FIG. 20, #144), and an address network coupling the active device and the memory subsystem (*See* e.g., FIG. 1 and FIG. 20, #150). The method includes a memory subsystem included in a node of the plurality of nodes receiving from an active device included in the node an address packet requesting an access right to a coherency unit. In response to receiving the address packet, the method includes the memory subsystem sending a responsive data packet to the active device dependent on a response indication associated with the coherency unit (*See*, e.g., FIG. 35 and specification page 103, lines 18-23). In addition, the method includes the node sending a coherency message requesting the access right to a different node of the plurality of nodes in response to a node identifier identifying the different node as a node in which the coherency unit is in a modified global access state (*See*, e.g., specification page 67, line 27-page 68, line 1).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claims 1-48 are rejected under 35 U.S.C. § 102(b) as being anticipated by Liencres et al. (U.S. Patent No. 5,434,993).

VII. ARGUMENT

First Ground of Rejection:

Claims 1-48 are rejected under 35 U.S.C. § 102(b) as being anticipated by Liencres. Appellant traverses this rejection for at least the following reasons.

Independent claims (by number):

Appellant respectfully submits that each of claims 1, 17, and 33 recites features not taught or disclosed by Liencres. For example, claim 1 recites features including: “each node includes and an address network and a data network respectively configured to convey address packets and data packets between the plurality of active devices and the memory subsystem;” and “wherein the node is configured to store a node identifier for the coherency unit, wherein the node identifier identifies a different

node of the plurality nodes in which the coherency unit is in a modified global access state.”

Appellant first notes that the Examiner, although allowed to interpret the art as broadly as is reasonable, has attempted to show that the system of Liencres, which includes nodes coupled to a common memory, but which have an inherently different structure, is the same as the system claimed by Appellant. Appellant respectfully disagrees with both the structure and function.

More particularly, on page 2 of the final Office action dated July 9, 2007, the Examiner responded to one of the Appellant’s arguments regarding whether element 33 of Liencres could be construed to be Appellant’s claimed address network. The Examiner asserts that it could because “The memory and processor are part of element 32 in Liencres, which is connected to element 33.” Appellant respectfully disagrees with the Examiner’s analogy. Specifically, as claim 1 explicitly recites “an address network and a data network, respectively configured to convey address packets and data packets between the plurality of active devices and the memory subsystem.” Thus, the address network and data network are separate entities. This is absolutely not disclosed nor is it inherent in Liencres. As shown below, Liencres merely discusses that packets are sent, but does not indicate how. In addition, the Examiner asserts element 33 is coupled to element 32 which has the memory (according to the Examiner) and the processor. However, claim 1 requires that the address and data networks convey address and data packets between the active devices and memory. If, as the Examiner suggests, element 33 is the address network, then it is NOT conveying address packets between the active device and memory.

The Examiner further rebuts Appellant’s arguments regarding whether the status bits of Liencres may be construed to indicate which node owns the coherency unit. The Examiner asserts that the status bits do identify the owner since they indicate/determine if the current node is the owner or not. Appellant submits that just because the status bits indicate the current node (i.e., the node within which the status bit resides) is or is not the owner, does not infer the status bits indicate the node (especially a different node) that is the owner.

Specifically, Liencres actually teaches at col. 7, lines 7-10

The bus cache controller 31 maintains a cache directory 46 containing the address tags and status bits for the data in the cache memory 37.
(Emphasis added)

Liencres also discloses at col. 1 lines 61-65

To maintain cache consistency, several status bits are usually maintained in the cache directory which reflects the current state of the information in each cache line. Common status bits maintained include a "valid" bit, a "shared" bit, and an "owned" bit. (Emphasis added)

Liencres further discloses at col. 8, lines 12-24

... The processor cache controller 35 issues a read request packet containing the required memory address to the bus cache controller 31 through the cache bus 33., the bus cache controller 31 responds to the read request packet by broadcasting a corresponding read request packet across the memory bus 25. The appropriate memory unit or processor subsystem on the memory bus 25 should eventually respond to the read request packet with a read reply packet containing the requested data. (Emphasis added)

Further, Liencres discloses in col. 7 “Read transactions”

When a memory request by the processor 21 cannot be fulfilled by the data in the processor cache memory 37, the processor cache controller 35 sends a read request packet across the cache bus 33 to the bus cache controller 31. The bus cache controller 31 proceeds to broadcast a corresponding read request packet across the memory bus 25. The read transaction initiated by the bus cache controller 31 consists of two packets: a read request packet sent by the bus cache controller 31 on the memory bus 25 and a read reply packet sent by another device on the memory bus. The read request packet contains the address of the memory requested by the processor cache controller 35 and is broadcast to all entities on the memory bus 25. A device on the memory bus 25 that contains the requested memory address responds to the read request packet with a read reply packet containing the subblock which includes the requested memory address. The read reply packet is generally issued by the main memory 23 except when the desired memory address is "owned" by another processor subsystem 20. In that case, the processor subsystem that owns the subblock must generate a read reply packet with the requested data.. (Emphasis added)

Thus in addition to the arguments above, from the foregoing disclosure, it appears Liencres is merely stating that the cache controller maintains status bits for each cache line and the status bits are what one would customarily expect (e.g., valid, shared, owned) in a conventional system. However, the status bits of Liencres do not indicate the different node in which the coherency unit is in a modified global access state. To the contrary, the status bits ONLY indicate whether the cache line is owned by the current node. In other words, the owned

bit in each node only indicates whether that node owns the coherency unit, and does NOT indicate the different node that does own the coherency unit. There is absolutely no teaching of any node identifier that identifies a different node of the plurality nodes in which the coherency unit is in a modified global access state.

Liencres is also disclosing in response to a read request, eventually the owner will respond with the data. Further, in the Read transactions section, Liencres is only teaching that read requests are met with replies that contain the data, and that the bus controller 31 broadcasts the request (because there is no indication that identifies which node).

In the rejection of claim 1, the Examiner asserts element 31 is the address network and the data network. Appellant disagrees for the same reasons that element 33 cannot be the claimed address and data networks. The bus controller 31 is not an address and data network respectively, meaning separate. Liencres is silent as to the actual configuration of the bus controller. All Liencres teaches is that the memory bus 25 and the cache bus 37 are packet buses. There is no teaching whatsoever that there are separate address and data networks.

Further in the rejection of claim 1, the Examiner points to Liencres col. 7, lines 7-10 to teach the node identifier. Appellant disagrees and as already discussed above, that particular passage in Liencres merely discloses status bits in the cache directory, which do not indicate the node in which the coherency unit is in a modified global access state.

Thus, Appellant submits Liencres does not teach or disclose “each node includes and an address network and a data network respectively configured to convey address packets and data packets between the plurality of active devices and the memory subsystem;” and “wherein the node is configured to store a node identifier for the coherency unit, wherein the node identifier identifies a different node of the plurality nodes in which the coherency unit is in a modified global access state,” as recited in claim 1.

For at least the above stated reasons, Appellant submits that the rejection of claims 1, 17, and 33 is in error and requests reversal of the rejection. The rejection of claims 2-5 and 7-16 (dependent from claim 1), claims 18-21 and 23-32 (dependent from

claim 17), claims 34-37 and 39-48 (dependent from claim 33) are similarly in error for at least the above stated reasons, and reversal of the rejection is requested. Each of claims 2-5, 7-16, 18-21, 23-32, 34-37, and 39-48 recite additional combinations of features not taught or suggested in the cited art.

Separately argued dependent claims (by number)

Claims 6, 22, and 38 depend from claims 1, 17, and 33, respectively. Accordingly, the rejection of claims 6, 22, and 38 is in error for at least the reasons highlighted above with regard to claims 1, 17, and 33. Additionally, each of claims 6, 22, and 38 recite features including: "... in response to receiving yet another coherency message acknowledging that the other one of the plurality of nodes gained the access right to the coherency unit, the interface included in the node is configured to update the node identifier to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit."

The Examiner asserts Liencres teaches this limitation in column 7 "write transactions." Appellant respectfully disagrees. More particularly, Liencres actually teaches

When a the cache memory system for a processor subsystem "owns" a particular cache line, it is allowed to modify the contents of the cache line. If the processor 21 modifies a cache line which is shared with other cache memories, the bus cache controller 31 performs a write transaction to update the information in the cache memories that share the cache line. If the cache line is subdivided into subblocks, not all modifications to data in a cache line result in a write transaction. In a system with subblocks it is only necessary to broadcast those subblocks which have been modified and reside in other caches as well. "Shared" flags are required for each subblock to keep this information, but will not be discussed here. The disclosure of U.S. patent application Ser. No. 07/620,496, filed Nov. 30, 1990, entitled "Consistency Protocols For Shared Memory Multiprocessors", now U.S. Pat. No. 5,265,235, issued Nov. 23, 1993, are incorporated by reference.
(Emphasis added)

Appellant fails to see how the above disclosure in Liencres teaches “update the node identifier to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit.” Again, Appellant submits the owned bit in each node only indicates whether each node owns the coherency unit, and does NOT indicate the different node that does own the coherency unit. Thus, the above passage cannot possibly teach updating the node identifier “to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit.”

For at least the above stated reasons, Appellant submits that the rejection of claims 6, 22, and 38 is in error and requests reversal of the rejection.

CONCLUSION

For the foregoing reasons, it is submitted that the Examiner's rejection of claims 1-48 is erroneous, and reversal of his decision is respectfully requested.

The Commissioner is authorized to charge any fees that may be due to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5681-13501/SJC.

Respectfully submitted,

/ Stephen J. Curran /

Stephen J. Curran
Reg. No. 50,664
AGENT FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: January 21, 2008

VIII. APPENDIX

The claims on appeal are as follows.

1. A system, comprising:

a plurality of nodes coupled by an inter-node network, wherein each node includes a plurality of active devices, a memory subsystem, and an address network and a data network respectively configured to convey address packets and data packets between the plurality of active devices and the memory subsystem;

wherein a memory subsystem included in a node of the plurality of nodes is configured to maintain a response indication indicating whether the memory subsystem should send a data packet corresponding to a coherency unit in response to receiving from an active device in the node an address packet requesting an access right to the coherency unit;

wherein the node is configured to store a node identifier for the coherency unit, wherein the node identifier identifies a different node of the plurality of nodes in which the coherency unit is in a modified global access state.

2. The system of claim 1, wherein each node includes an interface coupled to send and receive coherency messages on the inter-node network, wherein an interface included in the node is configured to store the node identifier for the coherency unit.

3. The system of claim 2, wherein the interface is configured to store the node identifier in a global information cache that includes storage for a plurality of node identifiers for a plurality of coherency units.

4. The system of claim 2, wherein the interface included in the node is further configured to store a global access state of the coherency unit in the node.
5. The system of claim 2, wherein in response to a coherency message sent by another one of the plurality of nodes requesting an access right to the coherency unit, the interface included in the node is configured to access the node identifier and to responsively send an additional coherency message to an interface included in the different node.
6. The system of claim 5, wherein in response to receiving yet another coherency message acknowledging that the other one of the plurality of nodes gained the access right to the coherency unit, the interface included in the node is configured to update the node identifier to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit.
7. The system of claim 6, wherein if the interface included in the node updates the node identifier, the interface is configured to send an address packet indicating a new value of the node identifier to the memory subsystem included in the node.
8. The system of claim 2, wherein the memory subsystem is configured to update the response indication in response to receiving address packets from active devices included in the node.
9. The system of claim 8, wherein the memory subsystem is configured to update the response indication to indicate that the memory subsystem should not respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting a write access right to the coherency unit from one of the active devices included in the node.

10. The system of claim 8, wherein the memory subsystem is configured to update the response indication to indicate that the memory subsystem should respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting to write a new value of the coherency unit to the memory subsystem.

11. The system of claim 8, wherein the memory subsystem is configured to not update the response indication in response to address packets requesting a shared access right to the coherency unit.

12. The system of claim 8, wherein the memory subsystem is configured to send a packet corresponding to the address packet to the interface included in the node if:

the response indication indicates that the memory should not respond;

the coherency unit is in a shared global access state in the node; and

the address packet requests write access to the coherency unit.

13. The system of claim 12, wherein in response to the packet corresponding to the address packet, the interface is configured to send a coherency message requesting write access to the coherency unit to the different node identified by the node identifier.

14. The system of claim 8, wherein the memory subsystem is configured to send a packet corresponding to the address packet to the interface included in the node if:

the response indication indicates that the memory should not respond; and

the coherency unit is in an invalid global access state in the node.

15. The system of claim 1, wherein the node is configured to update the node identifier in response to receiving coherency messages from other ones of the plurality of nodes via the inter-node network.
16. The system of claim 1, wherein an active device that sends the address packet specifying the coherency unit gains the access right to the coherency unit in response to receiving the data packet from the memory subsystem if the memory subsystem sends the data packet.
17. A node for use in a multi-node computer system, the node comprising:
 - a plurality of client devices including a memory subsystem, an active device, and an interface configured to send and receive coherency messages on an inter-node network coupling nodes in the multi-node computer system;
 - an address network configured to convey address packets between the plurality of client devices;
 - a data network configured to convey data packets between the plurality of client devices;
 - wherein the memory is configured to maintain a response indication indicating whether the memory should send a data packet corresponding to a coherency unit on the data network in response to receiving an address packet requesting an access right to the coherency unit from the active device;
 - wherein the node is configured to store a node identifier for the coherency unit, wherein the node identifier identifies a different node in the multi-node system in which the coherency unit is in a modified global access state.

18. The node of claim 17, wherein the interface is configured to store the node identifier for the coherency unit.
19. The node of claim 18, wherein the interface is configured to store the node identifier in a global information cache that includes storage for a plurality of node identifiers for a plurality of coherency units.
20. The node of claim 18, wherein the interface included in the node is further configured to store a global access state of the coherency unit in the node.
21. The node of claim 18, wherein in response to receiving a coherency message sent by another one of the nodes requesting an access right to the coherency unit, the interface is configured to access the node identifier and to responsively send an additional coherency message to an interface included in the different node.
22. The node of claim 21, wherein in response to receiving yet another coherency message acknowledging that the other one of the plurality of nodes gained the access right to the coherency unit, the interface is configured to update the node identifier to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit.
23. The node of claim 22, wherein if the interface updates the node identifier, the interface is configured to send an address packet indicating a new value of the node identifier to the memory subsystem.
24. The node of claim 17, wherein the memory subsystem is configured to update the response indication in response to receiving address packets from one or more active devices included in the node.

25. The node of claim 24, wherein the memory subsystem is configured to update the response indication to indicate that the memory subsystem should not respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting a write access right to the coherency unit from one of the one or more active devices.

26. The node of claim 24, wherein the memory subsystem is configured to update the response indication to indicate that the memory subsystem should respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting to write a new value of the coherency unit to the memory subsystem.

27. The node of claim 24, wherein the memory subsystem is configured to not update the response indication in response to address packets requesting a shared access right to the coherency unit.

28. The node of claim 24, wherein the memory subsystem is configured to send a packet corresponding to the address packet to the interface if:

the response indication indicates that the memory should not respond;

the coherency unit is in a shared global access state in the node; and

the address packet requests write access to the coherency unit.

29. The node of claim 28, wherein in response to the packet corresponding to the address packet, the interface is configured to send a coherency message requesting write access to the coherency unit to the different node identified by the node identifier.

30. The node of claim 24, wherein the memory subsystem is configured to send a packet corresponding to the address packet to the interface if:

the response indication indicates that the memory should not respond; and

the coherency unit is in an invalid global access state in the node.

31. The node of claim 17, wherein the node is configured to update the node identifier in response to the interface receiving coherency messages from other nodes in the multi-node system.

32. The node of claim 17, wherein the active device that sends the address packet specifying the coherency unit gains the access right to the coherency unit in response to receiving the data packet from the memory subsystem if the memory subsystem sends the data packet.

33. A method of operating a multi-node computer system, wherein the multi-node computer system includes a plurality of nodes coupled by an inter-node network, wherein each node includes an active device, a memory subsystem, and an address network coupling the active device and the memory subsystem, the method comprising:

a memory subsystem included in a node of the plurality of nodes receiving from an active device included in the node an address packet requesting an access right to a coherency unit;

in response to said receiving, the memory subsystem sending a responsive data packet to the active device dependent on response indication associated with the coherency unit;

the node sending a coherency message requesting the access right to a different node of the plurality of nodes in response to a node identifier identifying the different node as a node in which the coherency unit is in a modified global access state.

34. The method of claim 33, further comprising an interface included in the node storing the node identifier for the coherency unit and sending the coherency message to the different node.

35. The method of claim 34, wherein said storing comprises the interface storing the node identifier in a global information cache that includes storage for a plurality of node identifiers for a plurality of coherency units.

36. The method of claim 34, further comprising the interface storing a global access state of the coherency unit in the node.

37. The method of claim 34, further comprising:

in response to receiving a coherency message sent by another one of the nodes requesting an access right to the coherency unit, the interface sending an additional coherency message to an interface included in the different node identified by the node identifier.

38. The method of claim 37, further comprising:

in response to receiving yet another coherency message acknowledging that the other one of the plurality of nodes gained the access right to the coherency unit, the interface updating the node identifier to identify the other one of the plurality of nodes if an active device included in the other one of the plurality of nodes gained write access to the coherency unit.

39. The method of claim 38, further comprising the interface sending an address packet indicating a new value of the node identifier to the memory subsystem if the interface updates the node identifier.
40. The method of claim 34, further comprising the memory subsystem updating the response indication in response to receiving address packets from one or more active devices included in the node.
41. The method of claim 40, wherein said updating comprises updating the response indication to indicate that the memory subsystem should not respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting a write access right to the coherency unit from one of the one or more active devices.
42. The method of claim 40, wherein said updating comprises updating the response indication to indicate that the memory subsystem should respond to address packets requesting an access right to the coherency unit in response to receiving an address packet requesting to write a new value of the coherency unit to the memory subsystem.
43. The method of claim 40, further comprising not updating the response indication in response to address packets requesting a shared access right to the coherency unit.
44. The method of claim 40, further comprising the memory subsystem sending a packet corresponding to the address packet to the interface if:
 - the response indication indicates that the memory should not respond;
 - the coherency unit is in a shared global access state in the node; and

the address packet requests write access to the coherency unit.

45. The method of claim 44, further comprising the interface sending a coherency message requesting write access to the coherency unit to the different node identified by the node identifier in response to receiving the packet corresponding to the address packet.

46. The method of claim 40, further comprising the memory subsystem sending a packet corresponding to the address packet to the interface if:

the response indication indicates that the memory should not respond; and

the coherency unit is in an invalid global access state in the node.

47. The method of claim 33, further comprising the node updating the node identifier in response to the interface receiving coherency messages from other nodes in the multi-node system.

48. The method of claim 33, further comprising the active device that sends the address packet specifying the coherency unit gaining the access right to the coherency unit in response to receiving the data packet from the memory subsystem if the memory subsystem sends the data packet.

IX. EVIDENCE APPENDIX

No evidence submitted under 37 C.F.R. §§ 1.130, 1.131, or 1.132 or otherwise entered by the Examiner is relied upon in this appeal.

X. RELATED PROCEEDINGS APPENDIX

There are no related proceedings.